BEYOND THE SHADOWS

CCSLM: Walking Outside Plato's Cave of Neural Networks

"A large body of neural-network theory may have been studying the geometry of the shadow cast by context-dependent linear mappings, without identifying the underlying mechanism as the primary object."

IN THE CAVE: STUDYING THE SHADOWS

Neural-network theory has focused on the shadows—the piecewise-linear regions created by nonlinear activations and their complex geometries.

OUTSIDE THE CAVE: THE REAL OBJECTS

The world outside is context selecting linear mappings. The mappings are the computational primitives. Context chooses which one operates.

CONTEXT (CONTROL PATHWAY)

Independent of the data pathway. Examples: time, mode, task, position, robot state, sensor state, environment, user intent, etc.

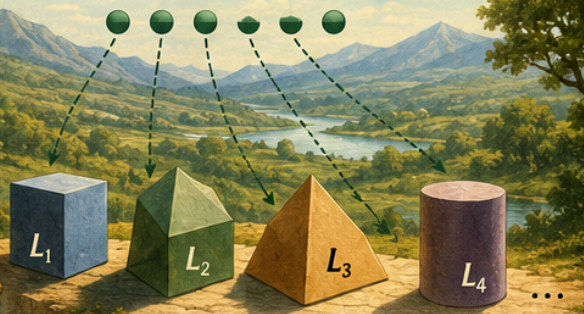

CONTEXT-SELECTED LINEAR MAPPINGS

A library of linear mappings. Context selects which mapping is applied at each layer.

THE SHADOWS

Observed behaviors, linear regions, decision boundaries, expressivity measures.

THE PRISONERS

Researchers building theory to explain the shadows they can see.

THE CHAINS

The assumption that computation must come from nonlinear neuron-by-neuron transformations.

CCSLM: THE MECHANISM MADE EXPLICIT

Computation is a sequence of context-selected linear mappings.

$$y = L_n(C_n, x_n) \cdots L_2(C_2, x_2)\, L_1(C_1, x_1)\, x$$

↑ output — ↑ context at layer i — ↑ data at layer i — ↑ input

THE TRADITIONAL VIEW (INSIDE THE CAVE)

Nonlinear activations create many linear regions.

$$y = W_n D_n(x_n) \cdots W_2 D_2(x_2)\, W_1 x$$

- Expressivity comes from activation patterns
- Study the number and geometry of linear regions
- Neurons as nonlinear feature detectors
- Complexity grows with depth and width
- Implicit switching via data-dependent nonlinearities

THE CCSLM VIEW (OUTSIDE THE CAVE)

Context selects the linear mapping directly.

$$y = L_n(C_n, x_n) \cdots L_2(C_2, x_2)\, L_1(C_1, x_1)\, x$$

- Expressivity comes from context-conditioned mappings
- Study the library of mappings and context structure
- Linear mappings as associative memories (with capacity)
- Complexity from context bandwidth and mapping diversity
- Explicit switching via an independent control pathway

WHY IT MATTERS

- Cleaner abstraction: Separate data from control.
- More efficient: Select, don't generate.
- Natural for real systems: Robots, time, modes, hysteresis, tasks, etc.
- Reinterprets decades of neural-network theory.

THE INVITATION

Neural networks may be the shadows. Context-selected linear mappings may be the objects. CCSLM proposes walking outside the cave.

SEE THE MECHANISM • MEASURE THE CONTEXT • SELECT THE MAPPING • UNDERSTAND THE COMPUTATION

2D Shadows 3D World.